
# PCIe® 4.0 Electrical Update

**Dan Froelich**

**EWG Co-Chair**

**Intel Corporation**

# Disclaimer

The information in this presentation refers to specifications still in the development process. This presentation reflects the current thinking of various PCI-SIG® workgroups, but all material is subject to change before the specifications are released.

# Outline

- **PCIe® 5.0 Rev 0.3**
- **PCIe® 4.0 Overview**
- **Lane Margining**
- **PCIe Channel Description**
- **SRIS**
- **Transmitter**
- **Receiver and Reference Clock**

# 5.0 Revision 0.3

- **The 5.0 signaling rate addition is 32 GT/s NRZ**
- **The maximum pad to pad loss target is expected to be around 35 dB**

- **Most areas in the electrical chapter are expected to change with the 32 GT/s signaling rate – the following areas are not expected to change**
  - Raw BER target remains E-12
  - # of Tx EQ taps and coefficient range
  - Tx EQ presets
  - High level channel compliance methodology with reference receiver and eye target
    - We may consider additional constraints such as insertion loss deviation and/or cross-talk magnitude
  - Backward compatibility with all slower speeds required
  - Tx voltage parameters
  - Tx and Rx stressed eye high level measurement methodologies (details will change)
    - Tx breakout channel de-embedded
    - Rx stressed eye methodology
  - AC coupling cap requirements
  - Transmitter and Receiver termination requirements
  - Reference clock voltage specifications with the possible exception of maximum slew rate

# PCIe 4.0 Overview

- **Key attributes of PCIe 4.0**
  - 16GT/s, using scrambling, same as 8GT/s
  - Maintains backward compatibility with installed base of PCIe devices
  - Limited channel reach: approx. 12” one connector (including 4” add-in card)
  - Longer channels require retimers or lower loss channels
- **New features**
  - Uniform spec methodology applied across all data rates (as possible)
  - Support for independent Refclk clocking mode with SSC (SRIS)
  - Integration of Retimer ECN
- **This presentation focuses on items adopted in the 0.9 specification as well as items being actively discussed in the EWG for adoption**
  - SRIS
  - Transmitter
  - Receiver
  - Channel

# Lane Margining

# Problem Statement (Lane Margining)

- **Rx CEM Compliance testing is expensive**
- **No standard way to test in L0**
- **No standard way to margin with silicon in production systems or in the field**
- **Retimers pose additional challenges**
- **Solution:**
  - Margin in L0 operation controlled by software from DSP and communicated by (enhanced) SKP OS (retimers) or registers (upstream ports)
  - Errors/ Status reported in CSRs by DSP and communicated by others to DSP using the (enhanced) SKP OS or registers (upstream ports)
  - Actual number of steps reported by each device (similar to TX EQ)
  - Add-in cards have to respond to commands – not required to drive margining on other devices
- **Also use the SKP OS mechanism to report errors in the "normal" (non-margin case) L0**

# Lane Margining Parameters (1)

| Parameter Name | Min | Max | Description |
|---|---|---|---|
| $M_{NumTimingSteps}$ | 6 | 63 | Number of time steps from default (to either left or right), range must be at least +/-0.2UI<br>Timing offset must increase monotonically<br>The number of steps in both positive (toward the end of the unit interval) and negative (toward the beginning of the unit interval) must be identical |
| $M_{MaxTimingOffset}$ | 20 | 50 | Offset from default at maximum step value as percentage of a nominal UI at 16.0 GT/s<br>A 0 value may be reported if the vendor chooses not to report the offset |
| $M_{NumVoltageSteps}$ | 32 | 127 | Number of voltage steps from default (either up or down), minimum range +/-50mV as measured by 16.0 GT/s reference equalizer<br>Voltage offset must increase monotonically<br>The number of steps in both positive and negative direction from the default sample location must be identical<br>This value is undefined if $M_{VoltageSupported}$ is 0b |
| $M_{MaxVoltageOffset}$ | 5 | 50 | Offset from default at maximum step value as percentage of one volt<br>A 0 value may be reported if the vendor chooses not to report the offset when $M_{VoltageSupported}$ is 1b<br>This value is undefined if $M_{VoltageSupported}$ is 0b |
| $M_{SamplingRateVoltage}$ | 0 | 63 | The ratio of bits tested to bits received during voltage margining. A value of 0 is a ratio of 1:64 (1 bit of every 64 bits received), and a value of 63 is a ratio of 64:64 (all bits received). |
| $M_{SamplingRateTiming}$ | 0 | 63 | The ratio of bits tested to bits received during timing margining. A value of 0 is a ratio of 1:64 (1 bit of every 64 bits received), and a value of 63 is a ratio of 64:64 (all bits received). |

# Lane Margining Parameters (2)

| Parameter Name | Min | Max | Description |
|---|---|---|---|
| $M_{VoltageSupported}$ | 0 | 1 | 1b indicates that voltage margining is supported |
| $M_{IndLeftRightTiming}$ | 0 | 1 | 1b indicates independent left/right timing margin supported |
| $M_{IndUpDownVoltage}$ | 0 | 1 | 1b independent up and down voltage margining supported |
| $M_{IndErrorSampler}$ | 0 | 1 | 1b Margining will not produce errors (change in the error rate) in data stream (ie. – error sampler is independent)<br>0b Margining may produce errors in the data stream |
| $M_{MaxLanes}$ | 0 | 31 | Maximum number of Lanes that can be margined at the same time. Recommended that this value is equal to the number of Lanes in the Link<br>Note: This value shall not exceed the number of Lanes in the Link |
| $M_{SampleReportingMethod}$ | 0 | 1 | Indicates whether sampling rates ($M_{SamplingRateVoltage}$ and $M_{SamplingRateTiming}$) are supported (1) or a sample count is supported (0). One of the two methods is supported by each device. |
| $M_{ErrorCount}$ | 0 | 63 | If $M_{IndErrorSampler}$ is 1b this is a count of the actual bit errors since margining started.<br>If $M_{IndErrorSampler}$ is 0b this is the actual count of the logical errors since margining started. See the Physical Layer chapter for the definition of what errors are counted.<br>111111b indicates >=63 errors |
| $M_{SampleCount}$ | 0 | 127 | Value = 3*log 2 (number of bits margined).<br>Where number of bits margined is a count of the actual number of bits tested during marging. The count resets on new margin command and stops when error count limit is reached or sample count saturates (also stops margining)<br>The count resets to zero when a new margin command is received. |

# Voltage Margin Example

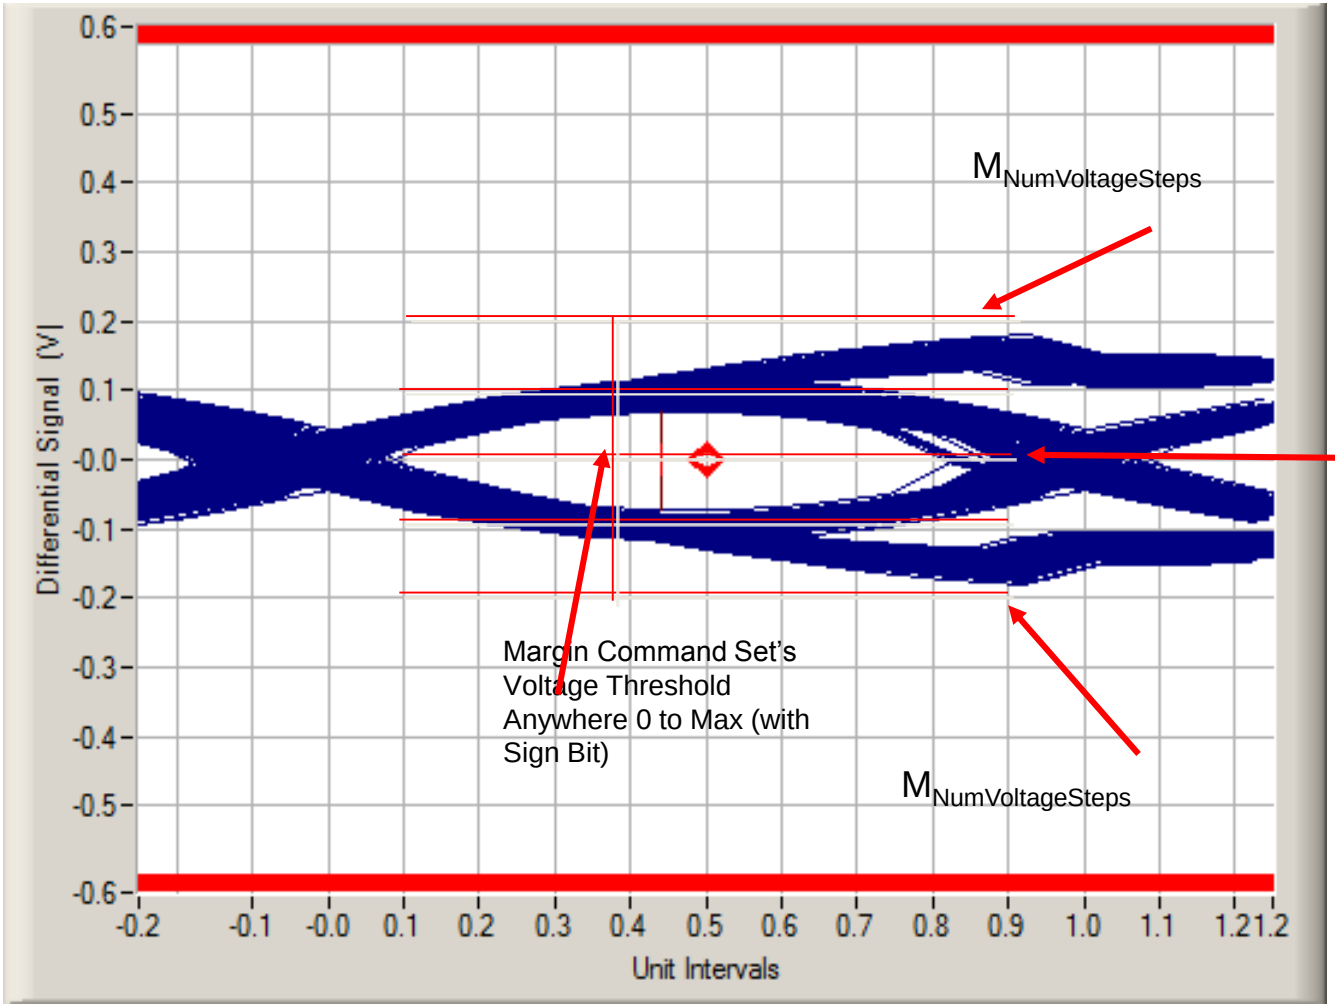

# Channel

# PCI Express® 4.0 Channels

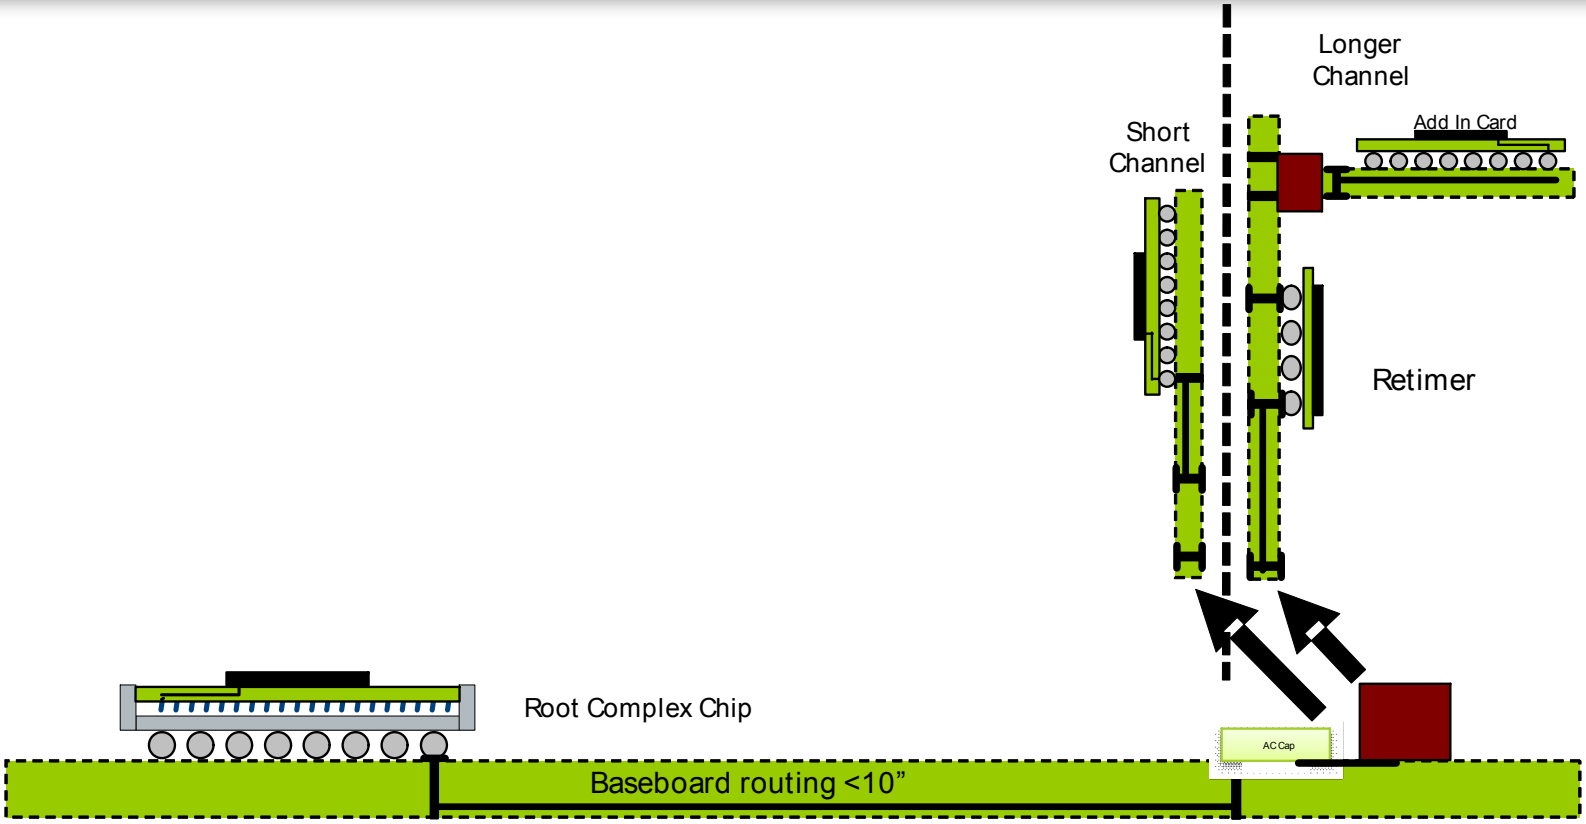


- End to end loss target ~= 28 dB
- Root Package loss ~= 5 dB
- Add-in Card Package loss ~= 3 dB
- Total Add-in Card ~= 8.0 dB
- Connector < 1 dB

# Minimum 4.0 Target Connector Performance

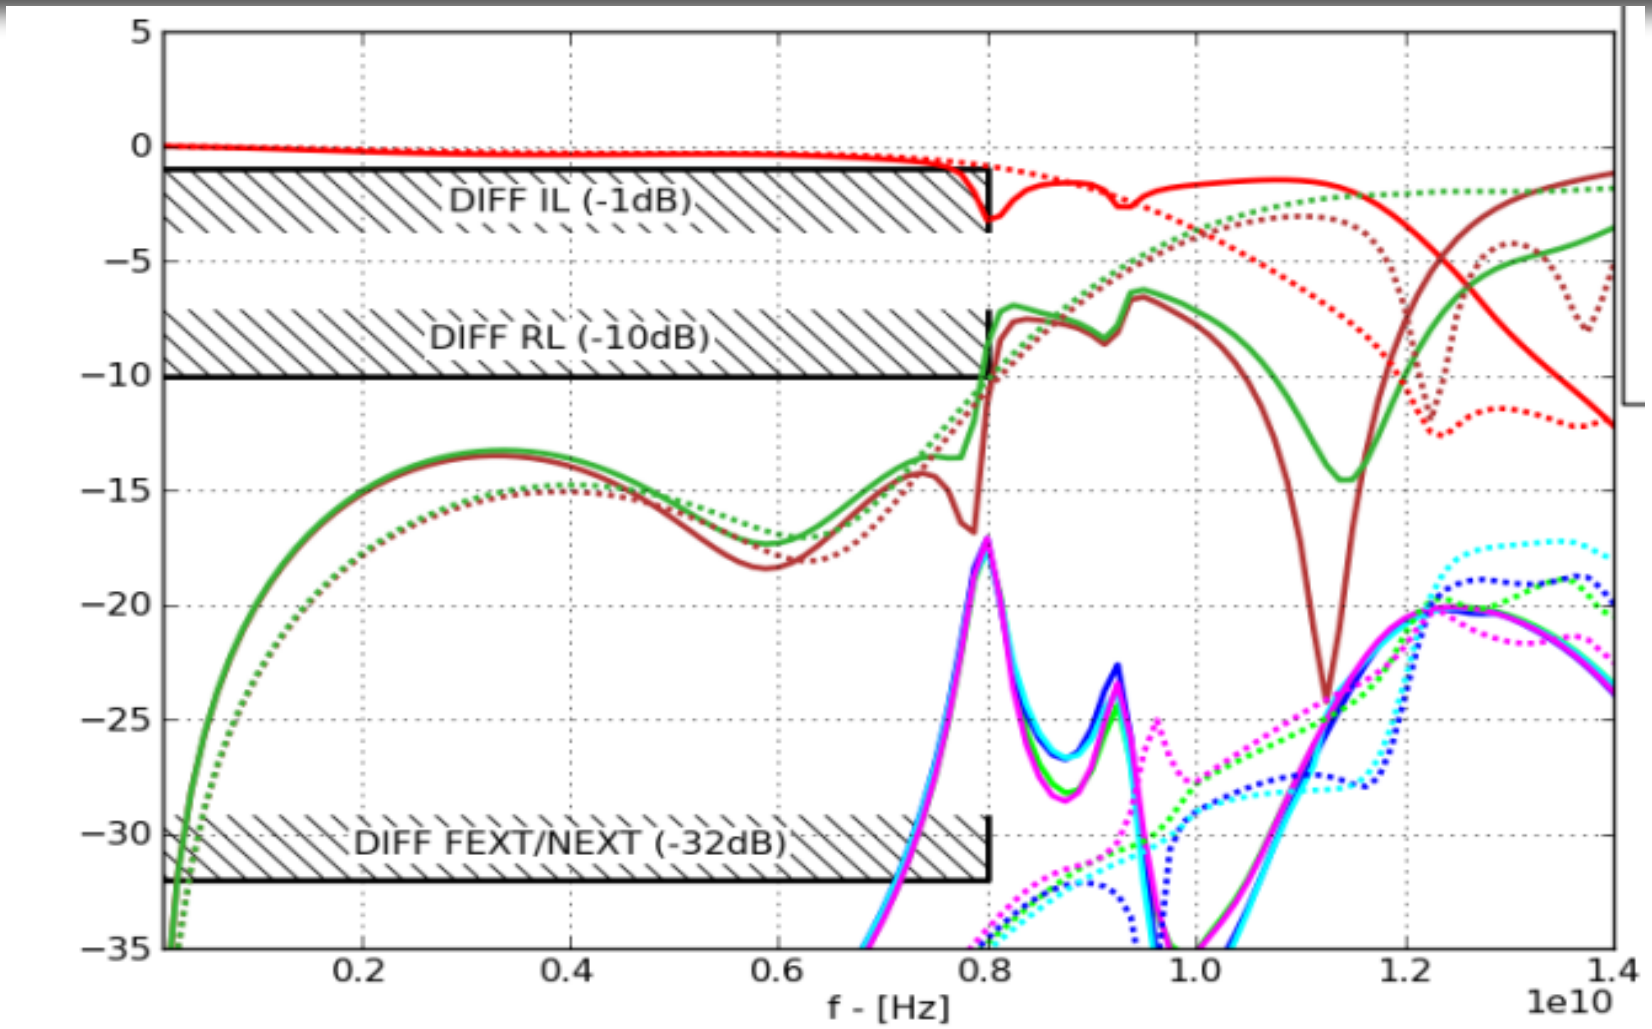


**Backwards Compatible Solutions Possible**

# Channel Loss Characteristics

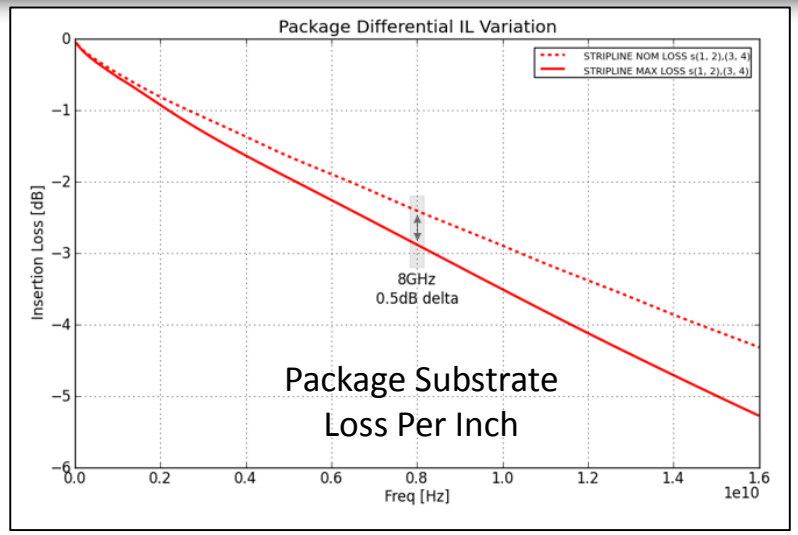


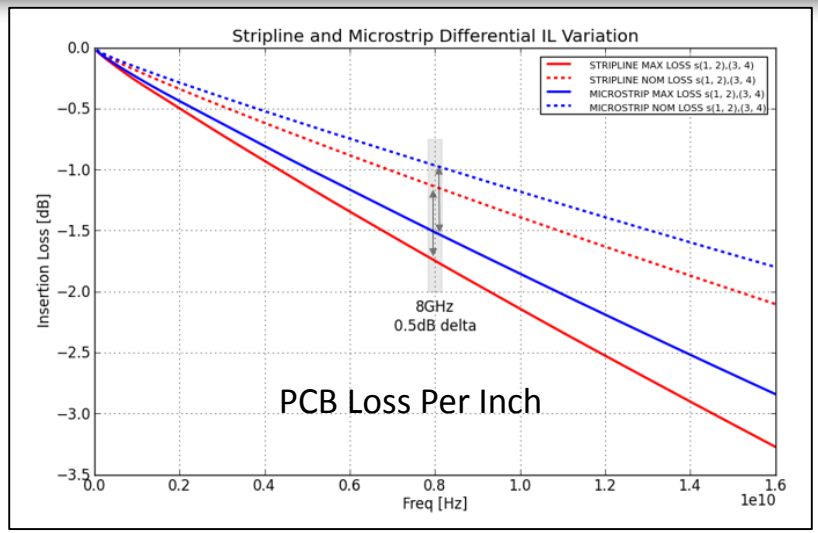


- **Package substrate & PCB loss per inch model used for channel study**
  - Tan delta [0.015 to 0.025]
  - Copper conductivity & roughness
  - Temperature & humidity variation
- **PCIE3 CEM cal channel RH & temp variation measurement @ 8GHz**
  - 12.5” Riser
  - 4” Mainboard
  - 2” CLB

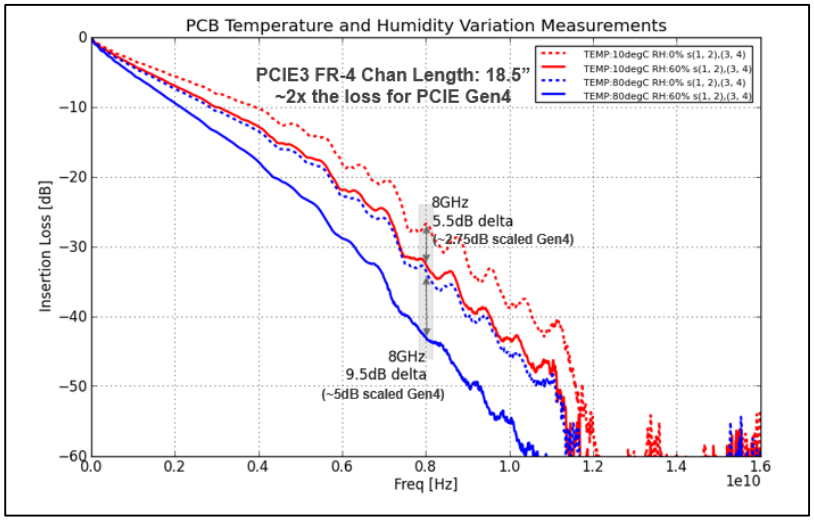

# CEM Channel Decisions

- **AIC Frequency Domain Loss Limit – 8 dB at 8 GHz**
  - An AIC ILpad-to-finger loss spec provides flexibility for implementation
  - Tradeoff package loss for PCB length, or pursue lower loss materials, etc.
  - Max trace length of 4” still possible with package improvement and/or lower loss materials

Loss budgets for channel max IL = -28dB:

| | Option 1 | | Option 2 | | Selected Limits | |
|---|---|---|---|---|---|---|
| Channel Max IL: -28dB | aic | mainboard | aic-1 | mainboard-1 | aic-2 | mainboard-2 |
| loss budget (dB) | 9.6 | 18.4 | 8.6 | 19.4 | 8.0 | 20.0 |
| pkg loss (dB) | 3 | 5 | 3 | 5 | 3 | 5 |
| via loss (dB) | 1 | 1 | 1 | 1 | 1 | 1 |
| CEM Connector (dB) | ---- | 1 | ---- | 1 | ---- | 1 |
| PCB Etch IL (dB) | 5.6 | 11.4 | 4.6 | 12.4 | 4.0 | 13.8 |
| etch Length @ -1.4dB/inch | 4 | 8.1 | 3.3 | 8.9 | 2.8 | 9.8 |
| etch Length @ -1.1dB/inch | 5.1 | 10.4 | 4.2 | 11.3 | 3.6 | 12.5 |
| etch Length @ -0.8dB/inch | 7.0 | 14.3 | 5.8 | 15.5 | 5.0 | 17.2 |

# Package Evaluation Process

- **Some packages are allowed to have insertion loss and/or cross-talk that exceed the reference packages**
- **Actual package performance must be evaluated**
- **Evaluation performed by running set of channel compliance simulations with reference channels replacing one reference package with package under evaluation**
- **Package with worse performance than reference must use actual package in channel compliance simulation and may use actual package in receiver calibration**

# SRIS

## (Separate Reference Clocks with Independent SSC)

# Inexpensive Cabling
# Independent Clock + Spread Spectrum (SSC) (SRIS)

- Challenge: PCI Express® (PCIe) spec did not support independent clock with SSC
  - SATA* cable ~ $0.50
  - PCIe cables include reference clock > $1 for equivalent cable
- PCIe base spec 3.0 ECNs approved
  - 1) Requires use of larger elasticity buffer
  - 2) Requires more frequent insertion of SKIP ordered set
  - 3) Requires receiver changes (CDR)
  - 4) Second ECN updates Model CDRs
- SRIS will create a number of new form factor opportunities for PCIe
  - OCuLink*
  - Lower cost external/internal cabled PCIe
  - Next generation of PCI-SIG* cable specification

*Example of Possible PCIe Cable*

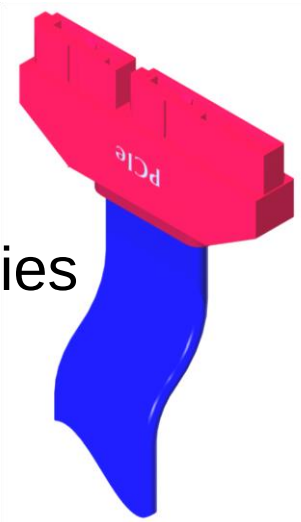

**Introduce New Terms for Separate Refclk Modes of Operation**
**5600ppm (New - SRIS) and 600ppm (Existing - SRNS)**

# Updated SRIS Model CDRs

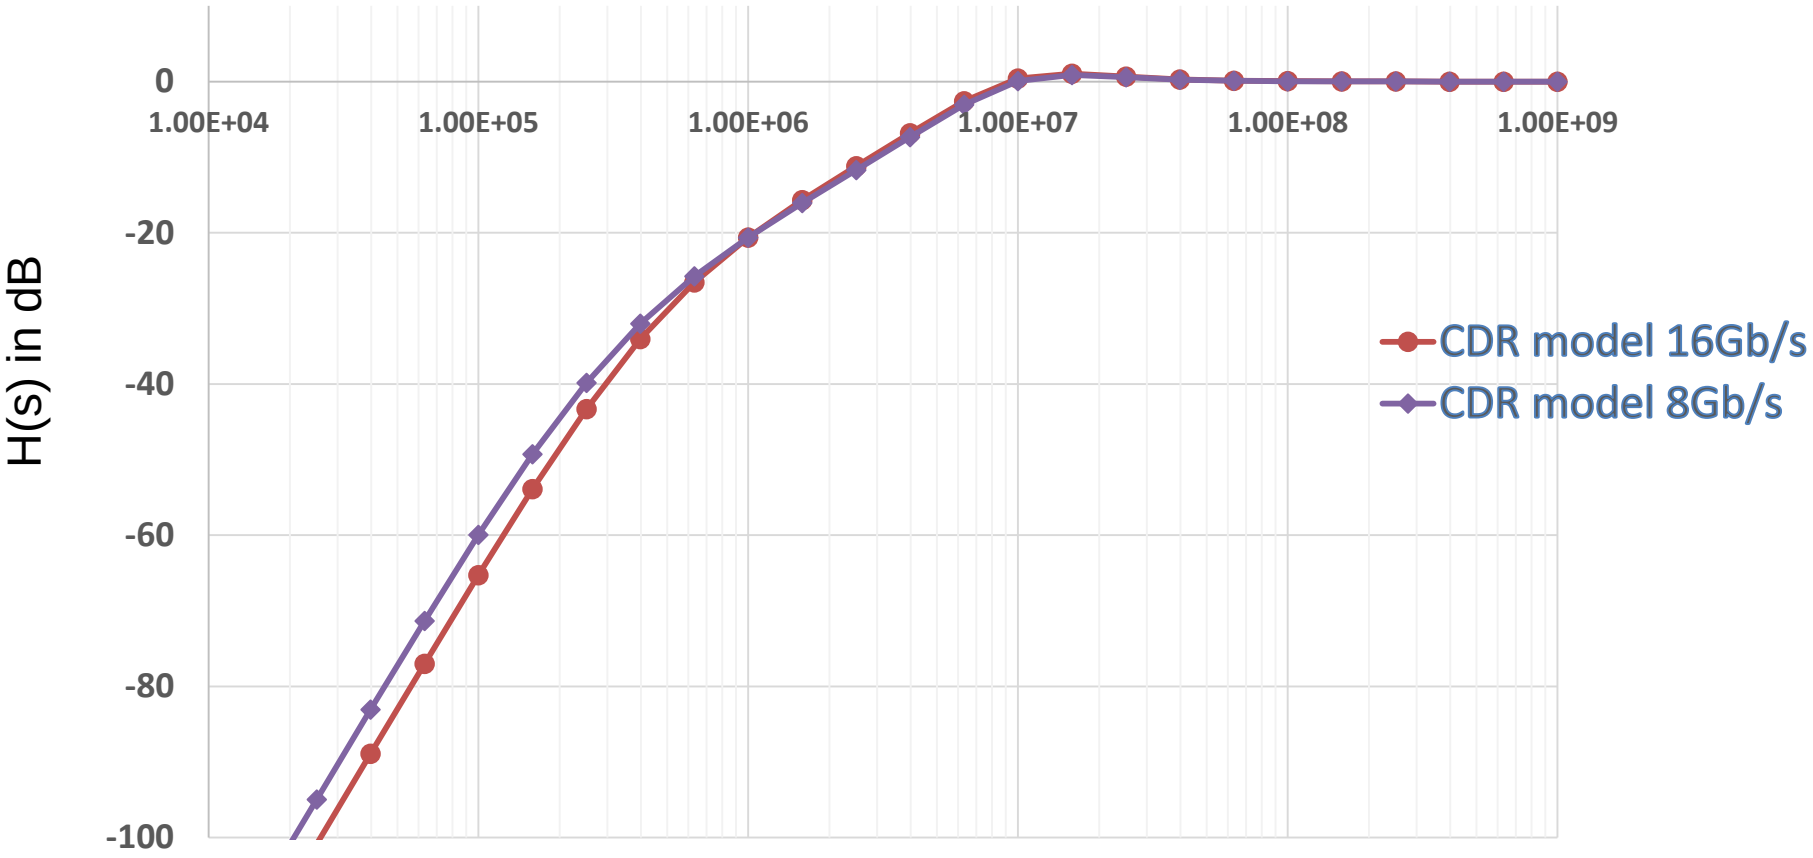


**More Rejection Was Needed at LF for Tx and Reference Clock Analysis**

# Transmitter

# Transmitter Specification

- **Preset definition**
  - Retain P0-P10 with same definition as PCIe 3.0 at 8GT/s
- **Package loss (ps21TX)**
  - Informative for root complex devices, informative for AIC devices
- **Architecture Specific Post Processing**
  - Common vs. Independent Refclk architectures
- **Jitter parameters**
  - Applied uniformly for all four data rates
  - Number of normative parameters reduced
  - Informative parameters added
- **Return Loss extended up to 8GHz**
  - T-coils likely required to meet limits
- **CTLE option for de-embedding breakout channel**

# Tx Testing: Base vs. CEM

- **CEM Tx Testing:**
  - End of reference channel waveform post-processed with reference equalizer. Best assessment of silicon-channel interaction.
  - Eye measured at a BER of 1E-12.
  - RJ & DJ gross separation performed at end of reference channel.
  - Eye may be closed at CEM connector for long channel designs.
  - Tx eye test only required to pass with optimal preset.

- **Mainboard Tx Testing:**
  - Non-ideal mainboard reference clock.
  - Suitable for testing commercially available mainboards.
  - Dual-port methodology samples Tx data lane and 100MHz ref clock.

# AIC Tx Package Loss Update

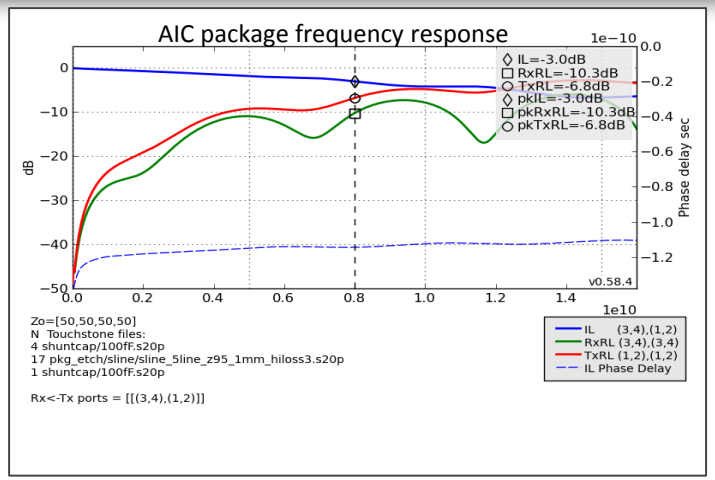


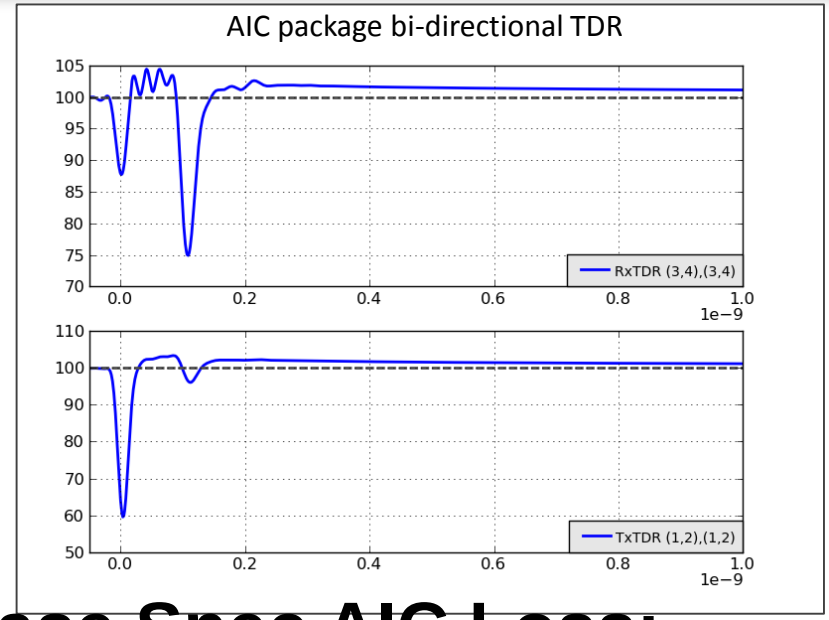


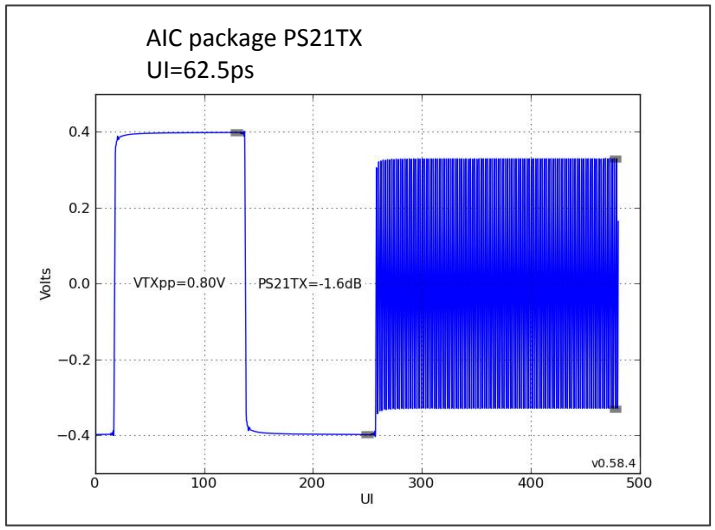


- **Base Spec AIC Loss:**
  - Informative IL limit = -3dB
    - Normative for Captive Channel Support
  - Tx-AIC-DEVICE PS21TX
  - Vendor must ensure Form Factor limits can be met

# EIEOS Update

- **Base spec requires a min of 175mVpk-pk at RX pin, using Preset10 (9.5dB boost +/-1.5dB)**
- **Rx die pad simulation using example worst-case -28dB server channel:**

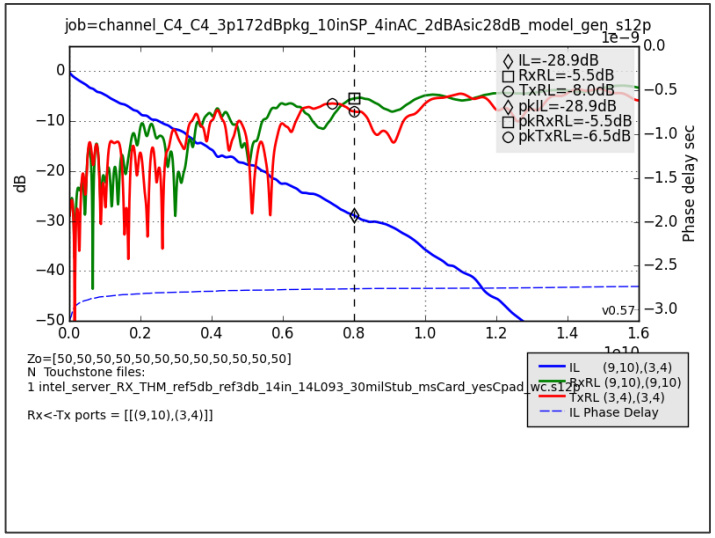


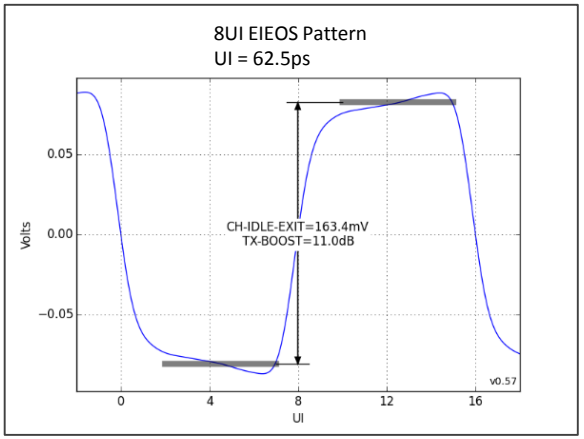


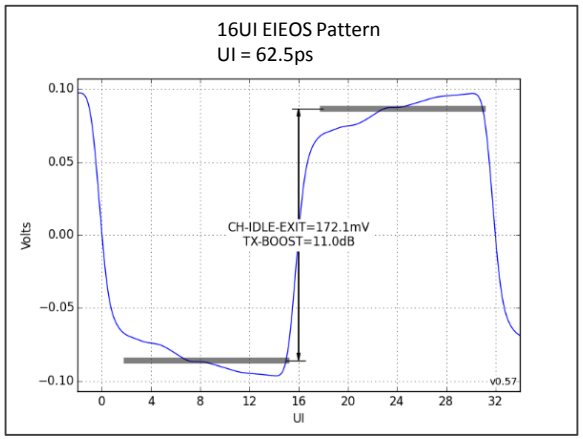


- **Electrical spec updated to use 16UI for EIEOS**

# Receiver

# Receiver Specification

- **Stressed eye methodology applied to all data rates**
  - Stressed jitter and voltage as a single test
- **Calibration channel defined by data rate dependent mask**
  - Direction to make loss variable at 16GT/s
  - Minimize Rj/Sj/DM variation across different set-ups
- **Separate Root Complex and AIC behav pkg models**
- **Behavioral Rx equalization data rate dependent**
  - 2.5 and 5.0G: none
  - 8.0 and 16.0G: CTLE and DFE (8G: 1 tap, 16G, 2 taps)
- **Eye height minimum reduced to 15mV for 16G**

# Receiver Linear Equalizer Pole 1 at 2 GHz

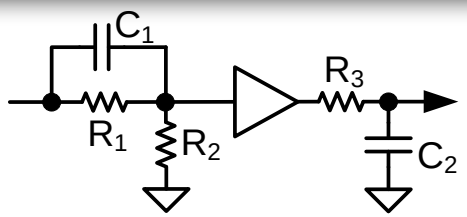


$$H(s) = \frac{sC_1R_1R_2 + R_2}{sC_1R_1R_2 + R_1 + R_2} \cdot \frac{1}{sC_2R_3 + 1}$$

$$H(s) = \frac{R_2}{R_1 + R_2} \cdot \frac{\frac{s}{\omega_Z} + 1}{\frac{s}{\omega_{P1}} + 1} \cdot \frac{1}{\frac{s}{\omega_{P2}} + 1}$$

$$G_{DC} = \frac{R_2}{R_1 + R_2}$$

$$\omega_{P1} = \frac{R_1 + R_2}{C_1R_1R_2}$$

$$\omega_Z = \frac{1}{C_1R_1} = G_{DC}\omega_{P1}$$

$$\omega_{P2} = \frac{1}{C_2R_3}$$

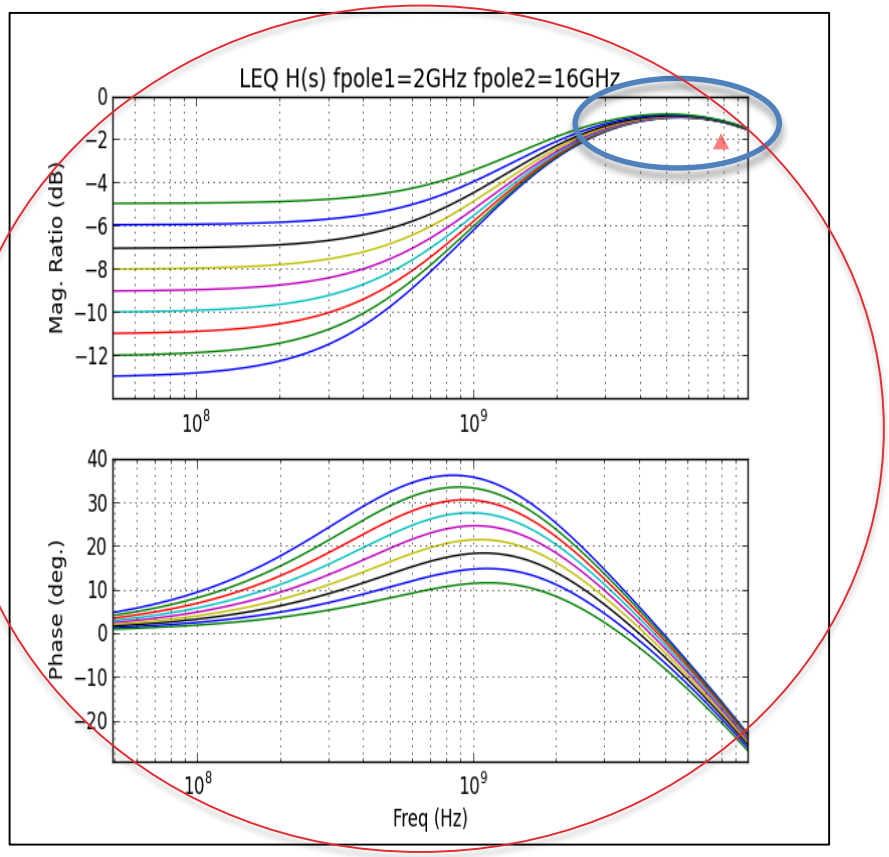


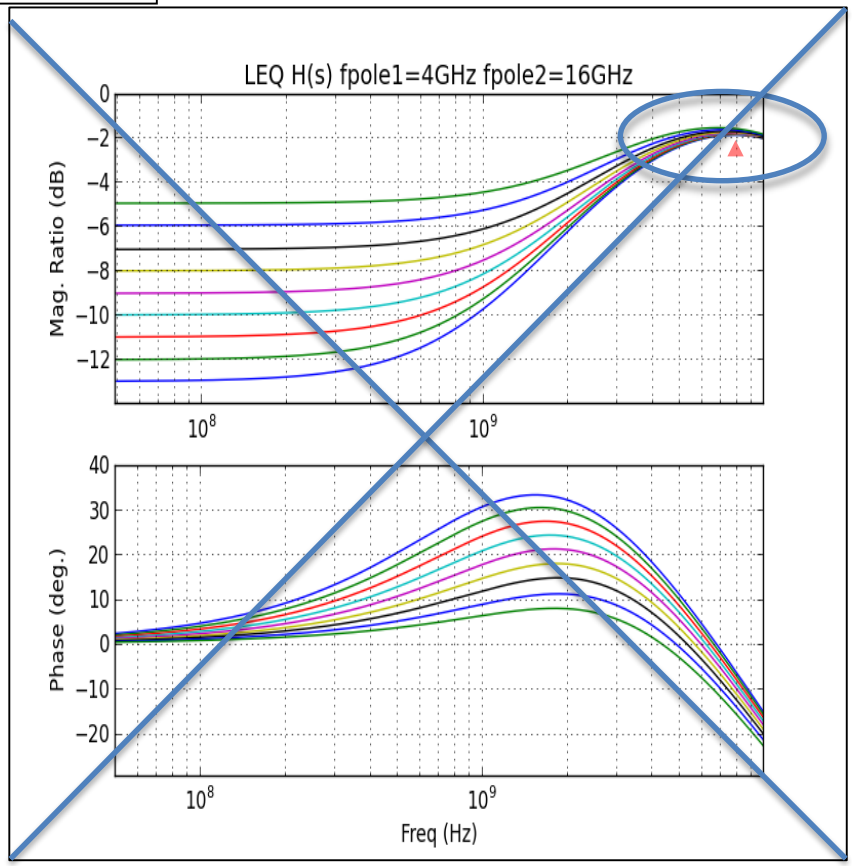

# Calibrating Stressed Eye

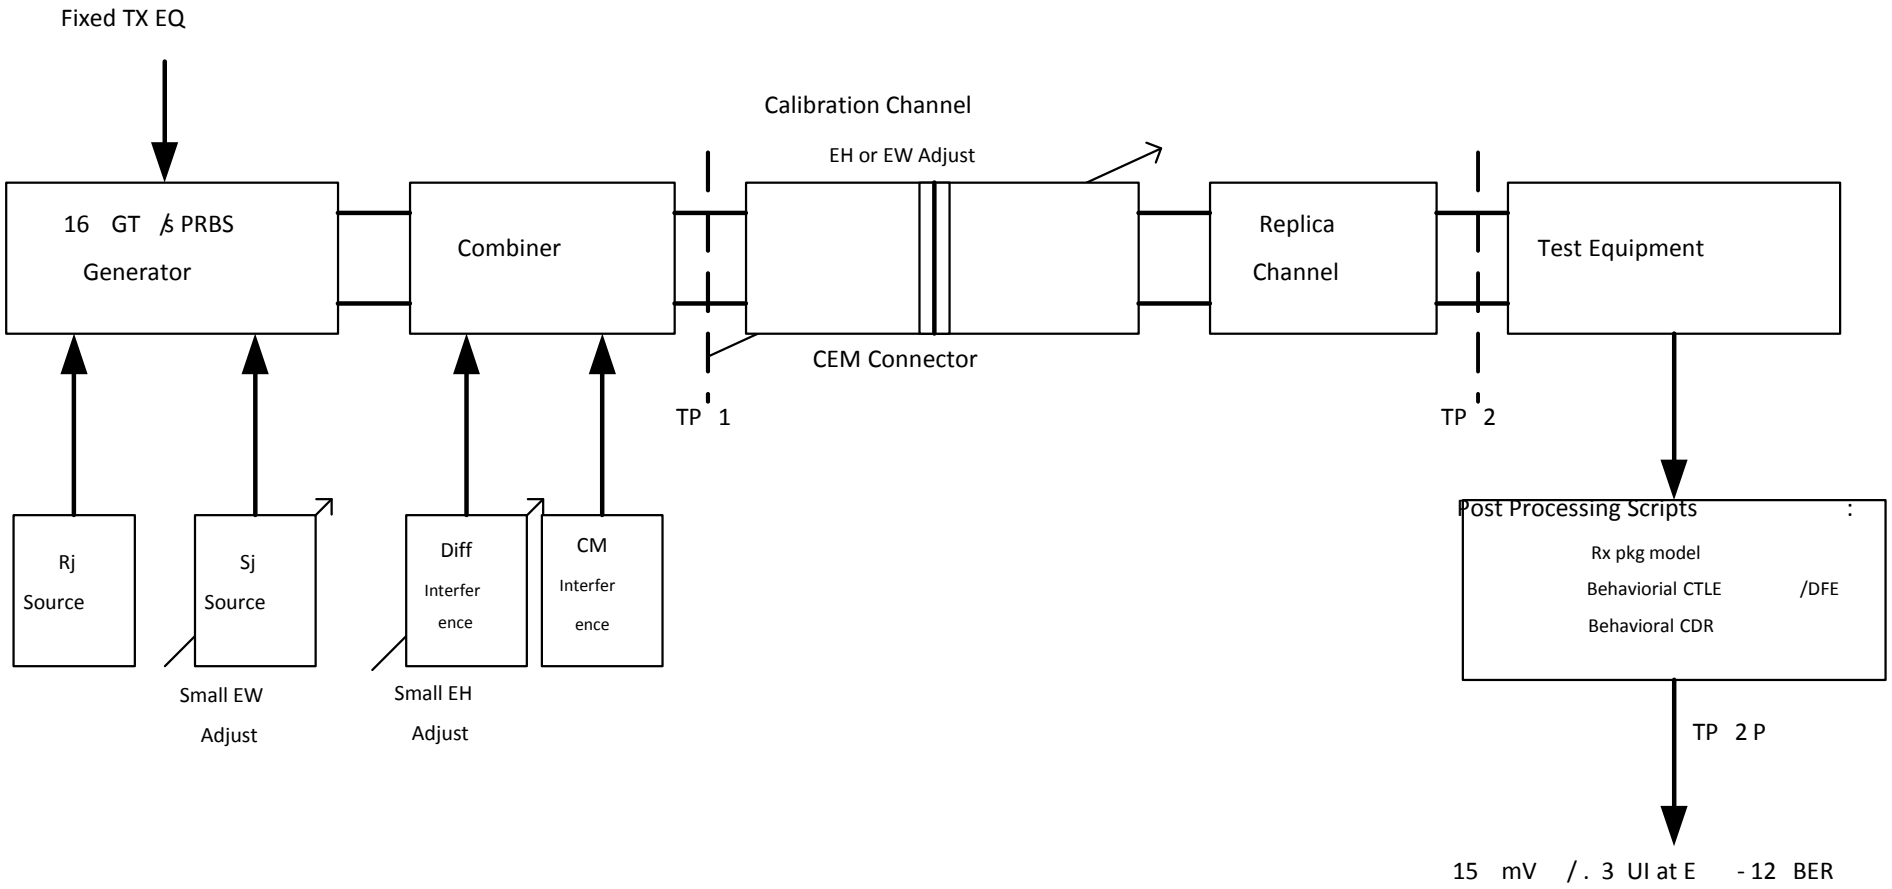

# Rx Cal Details
# Differential Noise and SSC Residual

| Symbol | Parameter | 2.5 GT/s | 5.0 GT/s | 8.0 GT/s | 16.0 GT/s | Units | Details |
|---|---|---|---|---|---|---|---|
| $V_{RX-LAUNCH}$ | Generator launch voltage | 800-1200 | 800-1200 | 800-1200 | 800-1200 | mV PP | Note 1 |
| $T_{RX-UI}$ | Unit Interval | 400 | 200 | 125 | 62.5 | ps | |
| $T_{RX-ST}$ | Eye width at TP2P | <0.4 | <0.32 | 0.30 | 0.30 | UI | Note 3,4, 10 |
| $V_{RX-ST}$ | Eye height at TP2P | 175 | 100 | 25 | 15 | mV PP | Note 2,4, 9 |
| $T_{RX-ST-SJ}$ | Swept Sj | N/A | N/A | Figure 9-29, Figure 9-30 | Figure 9-29, Figure 9-30 | UI PP | Note 5, Sec. 9.4.2.2.1 ff. |
| $T_{RX-ST-RJ}$ | Random Jitter | N/A | 3.4 (max) | 3.0 (max) | 1.0 (max) | ps RMS | Note 6,7 |
| $V_{RX-DIFF-INT}$ | Differential noise (2.1 GHz) | N/A | N/A | 14 | 14 | mV PP | Note 6,7 |
| $V_{RX-CM-INT}$ | Common mode noise | 150 | 150 | 150 | 150 | mV PP | Note 8 |
| $V_{SSC-RES}$ | SSC Residual | N/A | 75 | N/A | 500 | ps PP | Note 11 |

Note 11: Applied for CC testing only as triangular frequency modulation with a frequency between 30 and 33 KHz.

# Example Rx Calibration Fixtures

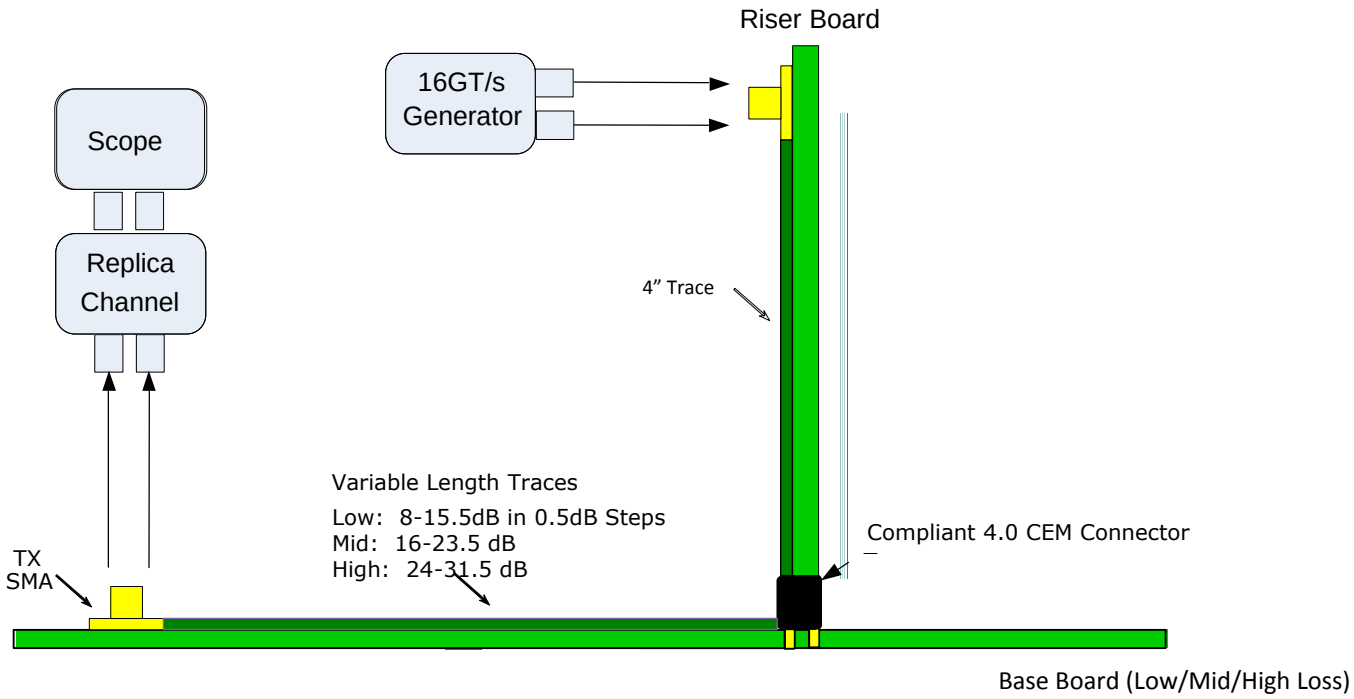

# Rx Cal Process (1)

- Calibrate the stress values to the nominal values in Rx Cal Details Tables
- Select an initial test channel length that give a loss at TP2P at 8 GHz of 27 dB +/- .5 dB.
- Measure the eye diagram for each TX EQ preset using the nominal TX Eq for the preset +/- .1 dB and select the TX EQ preset that gives the largest eye area.
- For all EH, EW and eye area measurements performed in receiver calibration the $A_{DC}$ in the reference receiver CTLE is varied over its minimum to maximum range in .25 dB steps. This is done to improve repeatability and accuracy in automated Rx calibration software and is only done for stressed eye calibration (not for channel compliance, etc.)
- Increase the calibration channel loss to the next available length/loss and measure the new eye diagram at the selected preset. Continue to increase the length/loss until either the height or width have fallen below the targets in Table 9-9 then the previous calibration channel length/loss is selected. If neither the height or width have fallen below the targets and the TP1 to TP2P loss at 8 GHz has reached 30.0 dB then advance to the next step.
- For the selected calibration channel length/loss, measure the eye diagram for each TX EQ preset and select the preset that gives the largest eye area. Note that this may be a different preset than step 3 due to the length/loss change.

# Rx Cal Process (2)

- **Adjust Sj, DM, and Voltage Swing to make final adjustments to the eye by sweeping them through the following ranges:**
  - Sj – 5 to 10 ps PP.
  - DM – 10 to 25 mV at TP2.
  - Differential Voltage Swing – 720 to 800 mV PP at TP1.
- **If the final Sj value is less then 0.1 UI then the Rj level is reduced so the eye width meets the target eye width with 0.1 UI of 100 MHz Sj.**
- **If there are multiple combinations of Sj, DM, and Voltage Swing that give valid solutions first pick the combination that is closest to the target eye width (18.75 ps). If there are multiple Sj, DM, and Voltage Swing combinations that are equally close to the target eye width then pick the one with Sj closest to nominal. The selected values must give a mean eye height and width (over at least 5 measurements – exact number of measurements needed for stable values will depend on lab set-up and tools) within the following ranges at BER E-12:**
  - Eye height – 15 mV +/- 1.5 mV
  - Eye width – 18.75 ps +/- 0.5 ps

# CC Reference Clock Specification (Implementation Trade-off for IR)

| Data Rate | CC jitter Limit | Notes |
|---|---|---|
| 2.5G | 108 ps pp | 1, 2 |
| 5.0G | 3.1 ps RMS | 1, 2 |
| 8.0G | 1.0 ps RMS | 1, 2 |
| 16G | 0.5 ps RMS<br>Note that .7 ps RMS is to be used in channel simulations to account for additional noise in a real system. | 1,2, 3 |

Note:

1. The Refclk jitter is measured after applying the filter function in Figure 9-45
2. Jitter measurements shall be made with a capture of at least 100,000 clock cycles captured by a real time oscilloscope with a sample rate of 20 GS/s or great. Broadband oscilloscope noise must be minimized in the measurement.
3. For the 16 GT/s CC measurement SSC spurs from the fundamental and harmonics are removed up to a cutoff frequency of 2 MHz taking care to minimize removal of any non-SSC content.

# Back-up

# EQ Tuning – High Loss Two Connector Server

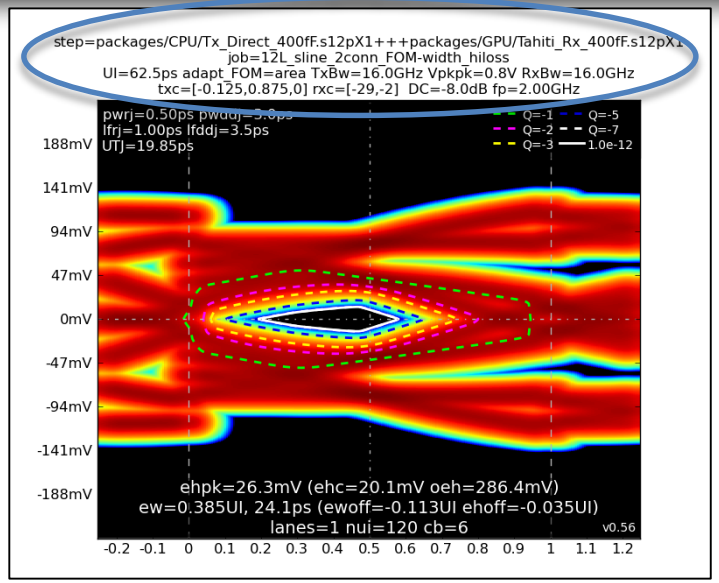


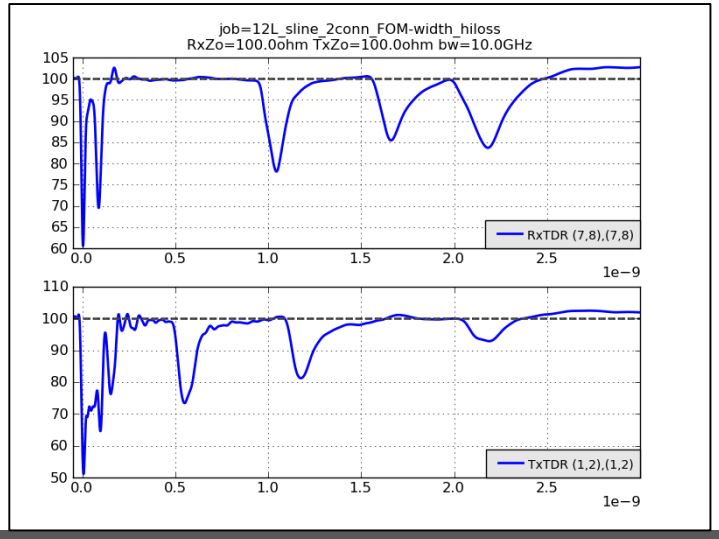


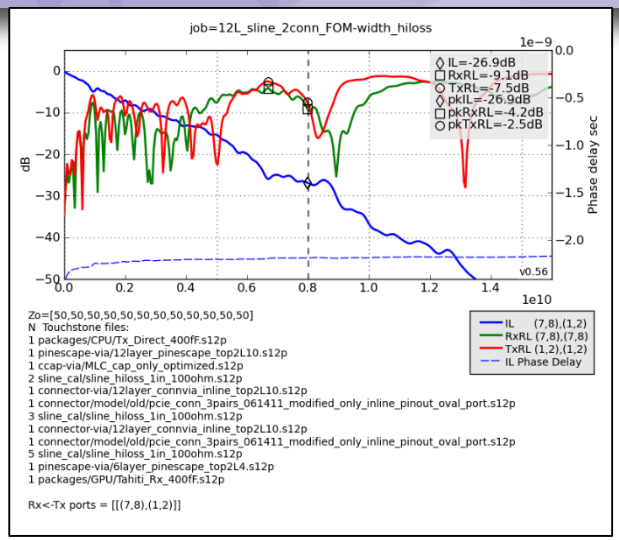


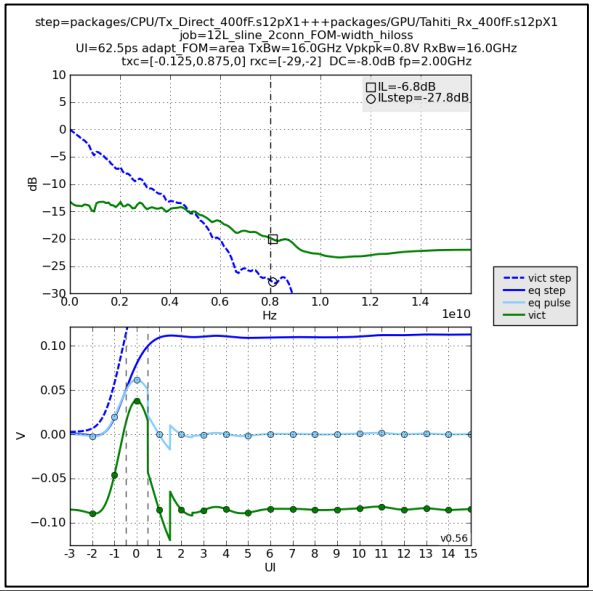

# EQ Tuning – Low Loss Two Connector Server

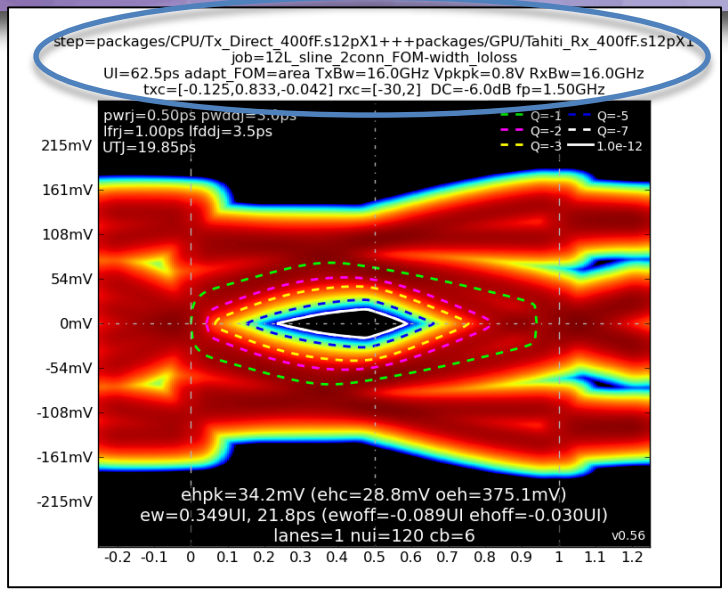


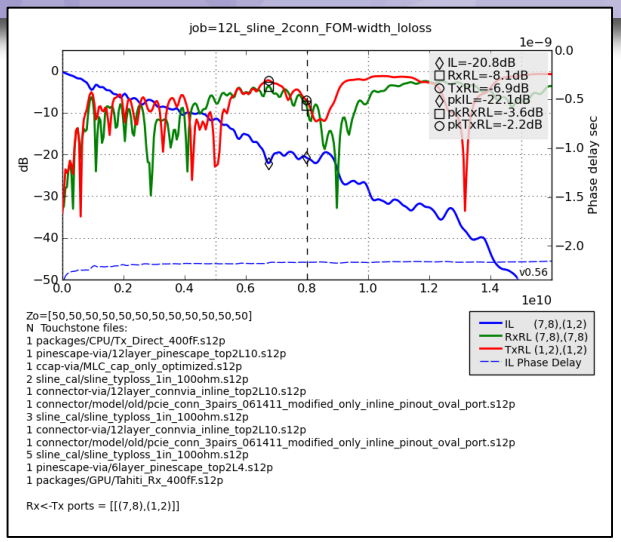


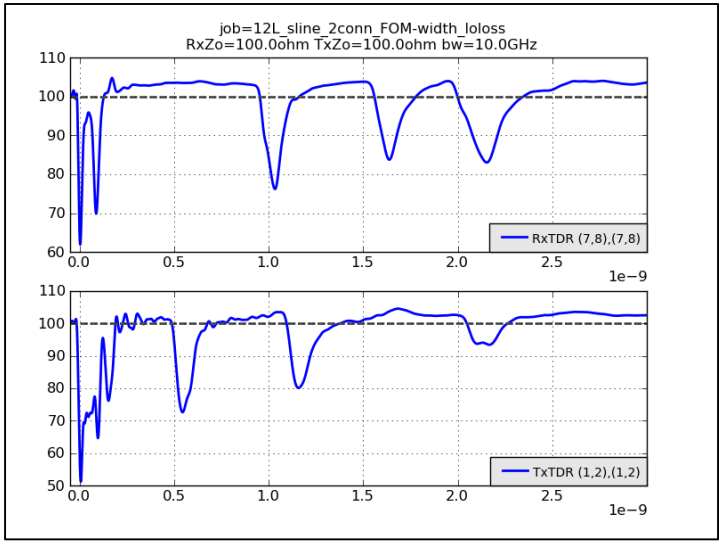


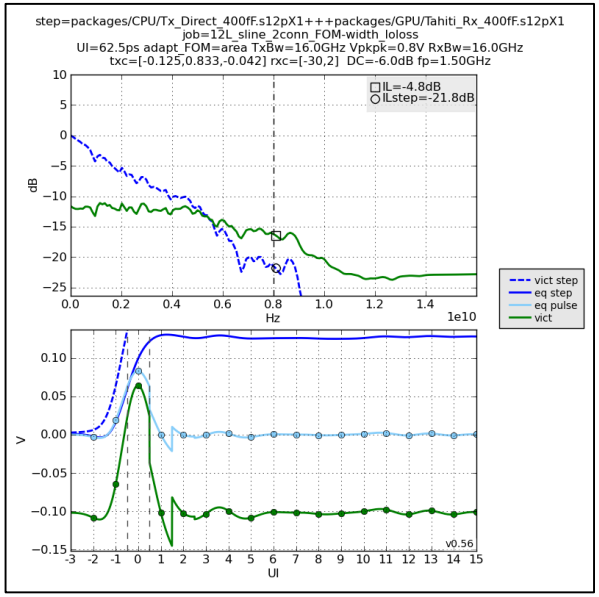

# Equalization Sweep

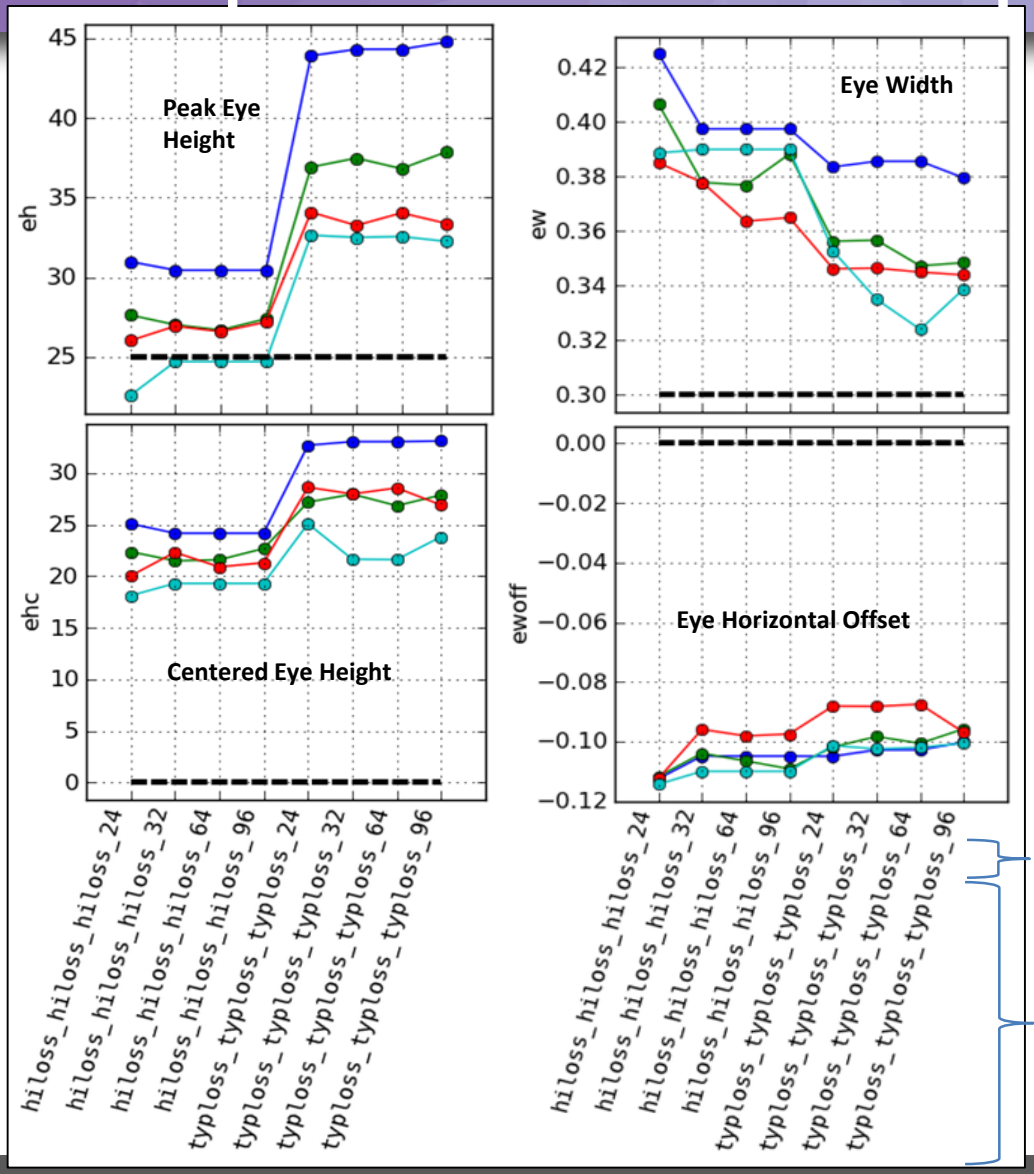


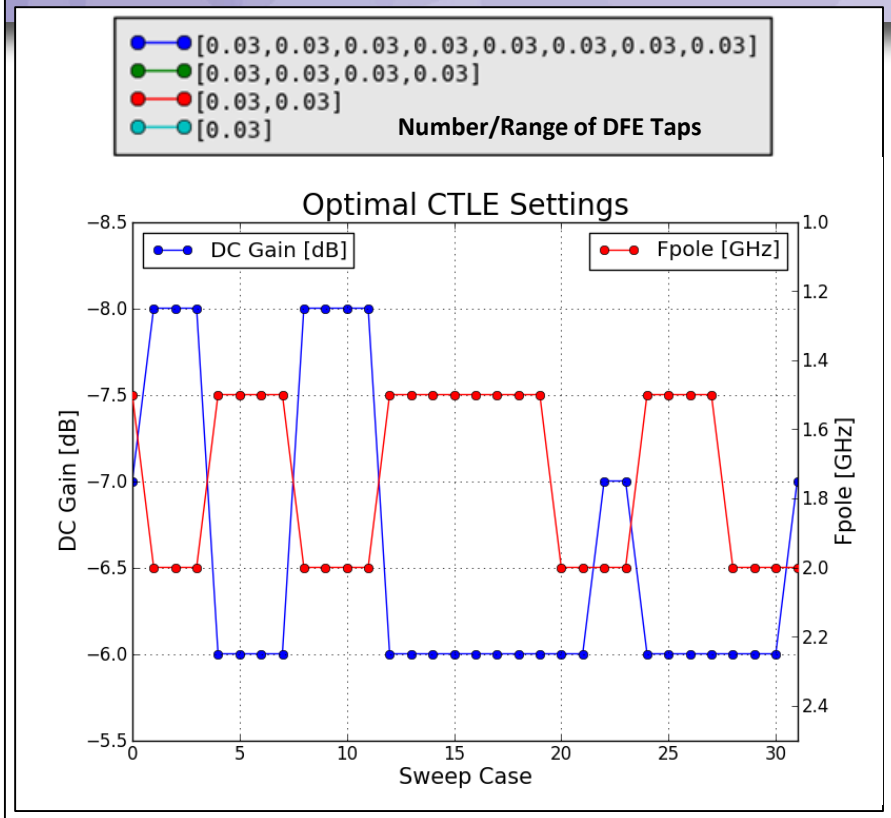


**Tx Tap Resolution**
**[1/24, 1/32, 1/64, 1/96]**

**Channel PCB Variation**
**[Tline1_Tline2_Tline3]**

# Waveform Post Processing Tool Requirements

- The test channel is the long Rx calibration channel with the Root reference package
- A step pattern with 256 ones and zeros is captured through the test channel by averaging 128 times on a real time oscilloscope for each preset.
- The channel compliance methodology is run and the TX EQ preset that produced the largest eye area is selected for exact eye height and width calibration.
- The channel compliance methodology is used with the selected TX EQ preset to produce an EH/EW of 15 mV and .3 UI @ E-12 BER by adjusting the Sj and voltage swing at the transmitter output.
- A pattern generator is calibrated to have the same jitter stress levels and Tx Swing as those used in the channel compliance simulations that produced an EH/EW of 15 mV and .3 UI @ E-12 BER at the pattern generator output.
- 2 million unit interval waveforms with compliance pattern are captured at each TX EQ preset at the end of the channel.
- For the preset that gives the largest eye area with the waveform post processing tool the EH and EW @ E-12 BER must match 15 mV and .3 UI within +/- 15%.

# Thank you for attending the PCI-SIG Developers Conference Israel 2017.

# For more information please go to www.pcisig.com